

[SQUEAKING] [RUSTLING] [CLICKING]

PROFESSOR: OK, so nice to see you all again. So we're going to talk about the family. So I'm going to today talk about basic intra-household bargaining models, sort of building-- in some sense kind of building us up to the Oudry reading that you guys did for today on sort of testing for Pareto efficiency in the household. And then Esther is going to come back and talk about other gender and within the family issues next time.

So just by way of background, right, so in some sense, so far, we've abstracted from households, right. So what I meant by that is, we assumed that a household had like a single utility function, u of x , where x is like everything that can be consumed. And we had some wealth or whatever. And we had the price of the pX is less than W , some budget constraint, right. So that's the unitary household model, OK.

But of course, as we know, the real world is more complicated than that, right. Households consist typically-- many households, their spouses, right. Often, people are married. And those spouses need to make decisions together about their consumption choices, OK. And that can be both statically and dynamically. I'm going to mostly talk about static choices today, but I'll mention, at least hopefully at the very end, an example of dynamic choice.

And also, you may imagine that different spouses may have different control over different types of assets or different consumption decisions. And this may be relevant for some development issues. So this is the benchmark. And the question is, is that kind of a useful benchmark, and when do we actually need to take into account the fact that actually there are different decision makers in the household, and they may be interacting in kind of different ways? Or when is a simplification kind of the good one?

OK, so there's sort of a long literature that started with these papers by Chiappori, Browning, and Bourguignon, which were about testing this unitary model and trying to see if we can accept it or reject it. And the paper you guys read for today is one example of one that's sort of about an example of testing the unitary model, but there are others.

So here's a model they're going to start with. They're going to start with a-- or here's how you build this up, right. So how would you build this up? You'd imagine, basically, that there's a-- I'll show you this in a second. We're going to have a joint utility function that's going to be some weighted sum of individual utility functions, right. And that can be very general. Like I can care about my spouse, or my spouse can care about me, or whatever. But it's individual utility functions.

And there are going to be some kind of weights between those two utility functions, OK. And those are going to be determined by bargaining or something like that. And what we're going to see is that we're going to see that basically, households are going to be maximizing this sum of weighted-- we're actually going to do sum of weighted utility. And that's going to imply a Pareto efficiency, right?

The Pareto efficiency means that if you can think about that there shouldn't be-- for example, as a test you guys saw in the Udry paper, the particular allocation of control over certain types of production in the household-- that shouldn't change the total amount produced in the household, for example. OK, but let's just show what this means, OK?

So imagine that basically, there's two household members, a man and a woman, just as an example, right? So the utility of each is going to be u_m and u_f , where q_m and q_f are the private consumption of the man and the woman. And Q is public consumption, OK?

So note, for example, this is very general, right? This can allow me to care about wife's consumption and vice versa, OK? But these are individual utility functions.

So we have some vector of consumption q and some vector of prices p . We also can allow the prices paid to be different, maybe. We set up some bargaining weights, λ -- by the way, if I'm going too fast, let me know. And also, the slides should all be posted on Canvas.

You should be able to have them. OK, so we saved some bargaining weights-- λ_m , λ_f . So the household utility is just given by the sum of the utilities, right?

So the household together is going to maximize some of these utilities subject to this joint budget constraint. And the joint budget constraint-- we're going to have the total consumption is equal to the total incomes. So note a couple of things that are important about this very simple setup, right?

One is the incomes just add together here on the right-hand side, right? There's nothing special about husband's income or wife's income. They're all just-- it goes into this collective income, right? And then they bargain over it, right?

But so this is going to be one of the tests that we're going to look for. That's going to be a test of income pooling, OK? The other thing that you should note is, tsh-tsh-tsh-tsh-tsh-tsh-tsh-tsh-tsh, this point, which is if we hold the bargaining weights fixed, then actually, even though this model had two different individuals who are making their decisions separately, this is actually going to look just like the unitary household again, right?

Why? Because if I hold the bargaining weights fixed, λ_m and λ_f , I can just define a new utility function, which is equal to the weighted sum of these utilities, right? And so this actually can be a challenge for testing these things in the data because they actually end up looking very similar.

They can look very similar in a lot of contexts, right? So even though here we might have different individual utilities and people care about different things, if the bargaining weights are fixed, then it ends up looking a lot like the unitary household. So what do we do with this?

OK, so basically, this is going to lead to two types of tests for whether or not we should be thinking of a unitary household decision-making. The first kind of test is going to come from this income pooling idea, which is to say, does income from different-- if, say, the man or the spouse, or one spouse or the other spouse gets income, does that affect the patterns of consumption, right? And the income pooling test says it shouldn't.

Second thing that we're going to come from is we're going to say, well, maybe we can think about something that affects the bargaining weights. And if I can think about different types of interventions, or policies, or things like that that are going to affect the λ_m 's and λ_f 's, and I say, well, if I change something that looks like it's going to increase the, say, the wife's bargaining power in a variety of ways, does that change the pattern of consumption? OK, and that's also going to be a test of this.

OK, so those are going to be two classes of tests. The third class of test is going to be the one in the Udry example, which is going to be a Pareto efficiency, which we'll talk about later. OK, so that's just the basic setup.

I went through that really fast, I think. But is that clear? Do you guys have questions? You should stop me if I'm going too fast. Good?

OK. So I think this is basically what I just said. So here are some examples of-- so I said we're not going to observe the bargaining weights.

We're going to make assumptions about them. And this is the income pooling test I already talked about. OK. Oh, the other thing which is going to be relevant for making these tests, for the empirical tests we're going to talk about, is we're going to have to go back to these q 's.

So we had q_f and q_m . Technically, if you just look at this model, the identity of what is a q_f and what is a q_m is irrelevant. Imagine I had a good exogenous shock in bargaining weights. If that changes the distribution of the q_f 's and the q_m 's, I can already reject a single unitary house model, right?

It doesn't matter what the different q 's are. But it might be nice to be a little more precise about our predictions, right? If we had certain ideas that there were, quote "women's goods" or quote "men's goods," then we might be able to generate a more specific test, a more directional test, which is to say, oh, if we have something that increases women in bargaining power and we know that women like such and such, we can look at these particular types of goods.

And I should say, actually, identifying that is hard, right? People make assumptions, sometimes about what are some goods that look like men's goods. But actually, there are many types of goods that we may have different preferences on, that men or women have different preferences on.

A lot of the tests will be like that. If you want to look at the number, the amount of goods that are strictly consumed by men or strictly consumed by women, for example, that can actually be pretty small, right? There are examples like that that people have used, but there's not a huge set that looks like that.

Plus, it also can be tricky if you have the altruistic utility, right? You might say, for example, maybe women in the household wear jewelry. And men tend not to wear jewelry. But of course, we all know that men can sometimes give jewelry to their wives, or vice versa. So it can be complicated.

But if we do have these more specific tests, sometimes, people will take the view of, oh, I can identify the following as goods that are associated with particular types of household members. And I can generate a directional test. If the income of that person increases, do we get more of their consumption?

That's a joint test of both income pooling and this different consumption. OK. So here are some examples, by the way, that people have used in the literature to think about some of these issues, right? And I'm curious.

Actually, maybe we can pause for a sec. So here are different things people have used as either sources of variation in the bargaining power of the household-- so different amounts of earned income from different household members, say maybe the age gap, resources brought to the marriage, unearned income, divorce laws. There's a whole literature-- actually, it started, I think, in the US on changing from making divorce easier as a change in bargaining power, where one is about the scarcity in the marriage market.

So what do you guys think of these as examples of things that you'd use to test for bargaining power? Do you have any thoughts on this? Seem sensible, not sensible? Are there other ones you'd put on this list?

First of all, do you understand why these things might be-- why you might think of-- so someone give me an example of how you might think of one of these as being related to bargaining power in the example I just set up. Somebody. Yeah? Remind me of name, sorry.

AUDIENCE: Whitney.

BEN OLKEN: Whitney.

AUDIENCE: So with the divorce law example, maybe you can think that if there's a law that makes it easier for women to get a divorce, then they'll have increased bargaining power in the household because they can say, if they don't like the way the things are going, they'll just leave.

BEN OLKEN: Right.

AUDIENCE: So they'll change the lambda that we have.

BEN OLKEN: Yeah, exactly. So if you think about-- how many of you guys have seen bargaining models in game theory, or in micro? Some of you. Only some of you.

Right. The way a typical bargaining model might work is to say we each get our outside option and some share of the surplus. That's the standard Nash bargaining solution, right? So imagine that basically, if divorce basically says my outside option has changed, if it's giving me a better outside option, if we don't agree on how we're going to share resources in this marriage, I can leave, right?

If that differentially increases the outside option of one side versus the other side, you could think of that as an increase in bargaining power, right? But it would have to be a divorce law that-- what's complicated is that you have think of it as-- from the model's perspective, if you're both stuck in this marriage, it's not totally obvious how moving to unilateral divorce necessarily increases-- it definitely means that both sides' outside options become more relevant.

But it's not obviously signed, I think. As an empirical matter, I think it does turn out to increase the woman's bargaining power. But I think from the model perspective, it's not totally obvious to me. Yeah, Wesley?

AUDIENCE: Can I ask a similar question about the age gap?

BEN OLKEN: Yeah.

AUDIENCE: It's not entirely clear to me which direction you would expect that to go in bargaining power. I could envision maybe you would expect that a larger age gap causes an increase in the bargaining power of a younger individual in some circumstances, maybe of the older individual in other circumstances. What's the take that I guess the age gap, or literature to try to [INAUDIBLE]?

BEN OLKEN: That's a good example. Actually, it's a good example. I'm trying to think of one off the top of my head.

You might imagine that, I think, in a bigger age gap between spouses-- particularly if spouses are getting married where one of them is very young, you might imagine that they're not as good bargainers, or as in a strong bargaining position. Their outside option may not be as good if they can't leave. Or you may just imagine that-- yeah, I think that's probably what I would say.

Other comments? Other thoughts? Yeah, Ahmed.

AUDIENCE: [INAUDIBLE] following up on that. We don't know the-- or can we say-- can we spot the age gap, for instance? Because I've heard stories about-- it can go either way. It can be--

BEN OLKEN: So give me an example of a story where it goes-- where a larger age gap-- so what do you mean?

AUDIENCE: I would assume, say, if it's a 30-year-old versus than 18-year-old, I would assume that age gap would be in favor of the 30-year-old. But if it's, say, a 30-year-old versus a 75-year-old, I would assume, at some point, it should go the other way.

BEN OLKEN: Yeah. So actually, so here's my basic takeaway from all this. The way I would want to think about all of this is to say, let's set up a bargaining model. And let's think about how these things affect your bargaining position.

A simple bargaining model where we each get our outside option plus some share of the surplus, right? That's not necessarily formally the same. But that's how I would think about changing this.

And I think the question is, how would you think about each of these? Which of these are really clear in that? And I agree with you, it's not totally obvious.

Maybe the 75-year-old, you're saying, has no outside option, right? But which of these do you think actually really does map very clearly into that threshold? Yeah, Patrick?

AUDIENCE: Earned income.

BEN OLKEN: Which one?

AUDIENCE: Earned income.

BEN OLKEN: Earned income? So say why.

AUDIENCE: You make the money, you call the shots.

BEN OLKEN: What? Sorry?

AUDIENCE: You make the money, you call the shots.

BEN OLKEN: Yeah, but try to put that into the bargaining perspective. So imagine-- how do you think about that from the bargaining perspective? So why is it the case that if you earn the money, you earn the shots?

AUDIENCE: OK. So if I try to--

BEN OLKEN: You call the shots.

AUDIENCE: --hold my money, or if I started bargaining with my money, I expect the utility out of this bargain model, having given in some of that into the pooling. It comes with a higher opportunity cost.

BEN OLKEN: Yeah. But let's try to be a little more precise, right? So imagine that basically, we have the-- let's call it the outside option of the man and the outside option of the woman. And we have, for example-- let's imagine we're going to have some, in this case, fixed bargaining weight alpha, right?

So what's going to happen is they're going to-- and if they can agree on everything together, they get some joint surplus s , right? So the man's outcome is going to be s minus om minus of times alpha plus om . And the woman's outcome is going to be 1 minus theta s minus om minus of plus of, right?

So this is the standard Nash bargaining solution with a constant, with simple case constant fixed bargaining weight alpha and a surplus from agreement that's just fixed or whatever, OK? So this is a simple setup I was describing. That basically says if we disagree and can't agree on anything, the man gets om , the woman gets of -- sorry, the man gets om , the woman gets of .

If we agree, then they each get their outside option-- om and of , respectively. And then they just share the surplus from agreement. It's s minus each of their outside options.

And they just share that alpha 1 minus alpha, OK? So that's the simple bargaining setup. So Patrick, going back to your example, how do you think about earned income in this example?

AUDIENCE: I would think that, for example, the [INAUDIBLE] much higher income.

BEN OLKEN: Sorry, you've got to speak up because of the mask. It's hard for me to hear you.

AUDIENCE: The male has a higher income [INAUDIBLE].

BEN OLKEN: Or one spouse, we can call it.

AUDIENCE: OK, one spouse has higher income. om is much larger. So if om is much larger, then the [INAUDIBLE] would be relatively smaller.

BEN OLKEN: Yeah. So by the way, I'm not being-- let me just say, I'm not being totally precise here. Because in the model I had over there, I had fixed bargaining weights over-- I had fixed bargaining weights over the utilities. And here I'm actually having fixed bargaining weights and they're just sharing the surplus. So this is not a precise mapping 1 to 1 .

And if I had thought this through more clearly in advance, I probably would have come up with a micro [INAUDIBLE] of that model over there with this one. But I didn't think it through in advance clearly enough. So my apologies.

But anyway, you can see, yes, exactly. So imagine you have more unearned income-- more earned income, om . What does that basically mean?

It means that in the event of the world in which we can't agree, I can go take my earned income and live a rich life. And if my spouse has a lower earned income, my spouse is going to go off and be poor. That basically means that I'm going to be able to extract more from the joint, more from the relationship.

And in this very simple model, this is not an income pooling model because the amount-- imagine that basically, this is the amount that the man, the utility that's going to be received by the man is going to depend on what fraction of the total income the man is bringing to the marriage and vice versa for the woman. Is that clear? OK. So I guess my point in having you think through this is when you see this stuff, I agree this one is, I think, pretty clear.

Unearned income also-- well, actually, no. Earned income is very clear, right? Because if we disagree, I can go off in the labor market and continue earning that income separately. So it's very, very easily mapped to my outside option. How do you think about unearned income? So first of all, what do we mean by unearned income?

AUDIENCE: Wealth?

BEN OLKEN: Huh?

AUDIENCE: Maybe wealth?

BEN OLKEN: Wealth, you said?

AUDIENCE: Yeah.

BEN OLKEN: But what do you mean?

AUDIENCE: Your estate.

BEN OLKEN: Huh?

AUDIENCE: Your estate.

BEN OLKEN: Your estate?

AUDIENCE: Yeah.

BEN OLKEN: Yeah, meaning what you inherit, for example. Yeah. So that's right. So OK. Someday, my computer will come back. Uh-oh. So what do you mean by--

AUDIENCE: So I guess it would depend on whether you've got a prenup, whether your unearned income is not going to be split between the two of you when you have that-- there's divorce.

BEN OLKEN: Exactly, right? So divorce law, for example, differs. In the US, for example, in some states-- and this certainly also differs by country-- we have what's called community property, which basically means everything is shared equally regardless of where it came from. But in other places, they keep track and say, if you were to divorce, you would get the-- but I came to the marriage with a lot more than my spouse did.

Then I would get more-- I would get that back. So I think this unearned income-- you have to think of it-- if you want to use it from this perspective, you have to understand what would actually happen in the event of a divorce, how that would affect your bargaining power, right? And it's not obvious it would-- if you're in a world where the states splits stuff up equally, then it's necessarily clear. It wouldn't matter.

So we're going to look at a paper in a second that looks at unearned income as opposed to earned income. Why might you prefer to do unearned income rather than earned income? If the empirical test is clearer for earned income, why do you think these papers might look at unearned income instead? Yeah.

AUDIENCE: [INAUDIBLE] from the [INAUDIBLE] on somebody?

BEN OLKEN: OK, exactly, right? So certainly, you can drop a cash transfer. And if I get to all the slides here, I have at least one paper on dropping cash transfers.

So yes, exactly. But suppose it was wealth, right? It was the wealth you bring to the marriage. Yeah, Wes?

AUDIENCE: Earned income assumes that everyone's in the labor market.

BEN OLKEN: What, sorry? Earning

AUDIENCE: Earned income would assume that both parties are working. In order to earn some income--

BEN OLKEN: Yeah, so right. So this is actually right. So as I think you were saying, this is totally endogenous. Earned income is endogenous.

What you actually need is a measure of your earning power, right? So for example, I could have great earning power as an MIT professor but be choosing not to work, for example. What we actually want to know is, what's my outside option if I were to go back into the labor market? What would my earned income be, right?

So you don't want actual earned income. But you want something proxy for earned income. For unearned income, what's the problem with this stuff, with resources brought to the marriage, for example? What's an empirical challenge, if you wanted to look at that?

AUDIENCE: You might not have data at all.

BEN OLKEN: Well, that's definitely true.

[LAUGHTER]

But suppose you did. So it seems exogenous, right? I don't control what I inherit. That's true.

AUDIENCE: [INAUDIBLE] to come dictate who you [INAUDIBLE].

BEN OLKEN: Absolutely, right? So exactly. So we're not going to talk a lot about marriage markets with prices. But certainly, there are certainly lots of marriage markets in the world, particularly in the developing world.

Explicitly, there are dowries and bride prices, right? And so there's an explicit marriage market that looks like that. Even without that, the marriage market, the matching market is pretty complicated.

And people may have different preferences in terms of who they're going to match with. And those things may be related to things that look like your determinants, look like your consumption preferences. Yeah.

AUDIENCE: So the concern here would be, for instance, that if you had, for instance, a disability of some sort that meant that you couldn't join the labor market, your parents might give you a higher dowry?

BEN OLKEN: No, that's not what I'm thinking. I'm thinking about-- I'm thinking of a much simpler thing. Sorry, that could be certainly true. Yes, that certainly could be.

I was thinking more-- who are the types of people who might, say, want to marry a really rich spouse, right? And who are people who don't care so much about marrying a rich spouse? They may have different consumption preferences.

So if I want to estimate a model where I'm going to-- what's the test, right? The empirical test was going to be something like-- I'm going to regress, say-- this is not really on the board. But I'm going to regress something about the q's of the-- I'll go back here so we're on the board.

But I'm going to regress, say, the man's consumption versus women's consumption as a function of how much bargaining power they have, or how much unearned-- what fraction-- for example, I'm going to run a regression of things that look like women's consumption on how much resources the woman brought to the marriage-- well, for the man, for example. But imagine that people's preferences over consumption are correlated with their preferences for who they're going to match with, right?

I want a really fancy life. Maybe I marry a really rich partner, for example, right? That could be an example where the endogenous composition of the marriage could be related to the preferences. You see what I'm saying?

AUDIENCE: Isn't that just for earned [INAUDIBLE]?

BEN OLKEN: Yes. Yeah. Yes, yes, it is. It's not necessarily better for unearned income. But yes, it's problem for earned income, too. Yeah.

AUDIENCE: So then a totally separate channel, but also a potential concern would be if your unearned income is somehow scaling inversely with your earning potential?

BEN OLKEN: Yes, that also could be a problem, too.

AUDIENCE: It would be a bad problem.

BEN OLKEN: Yes, yes. Correct. If those things were correlated too, that could be a problem.

So I'm going to show you a couple of examples of some specific papers. We can talk about this in specific context. Actually, the main point I wanted you to take away from this-- I'm putting this on the board-- is to start thinking about it from this perspective.

When we say these things impact bargaining power, what do we mean? The simplest version of that is what am I going to get in the outside option? And so whenever you see something like that, you should be thinking through, what is the outside option? How do these things map to the outside option? That was actually my main point. OK. Yeah.

AUDIENCE: Sorry, how should we think of outside in this model? You said it was--

BEN OLKEN: So I said it was exogenous.

AUDIENCE: Yeah, but--

BEN OLKEN: Where does it come from?

AUDIENCE: Yeah.

BEN OLKEN: So I don't have a good theory of alpha in this model. In some bargaining models, it comes from your discount rate. So this is the standard Nash bargaining model.

Some models microfound this with a series of alternating offers. So I make you an offer. And you can choose to accept it or make a counteroffer. Or we go back and forth and back and forth till we agree.

And in those models, I think that, if I remember correctly, the differences in the discount rates between the two parties are going to determine the alpha. That may not have been a clear-- that may not have been the answer you're looking for. But we treat it as-- from the Nash perspective, we treat it as exogenous.

And I think there are some models that microfound it. And I know that's an example of-- I think that's the Rubenstein model, I believe. OK, someone is nodding.

So OK. I already said that. OK, so here's the first paper I wanted to talk about in a bit of detail. This is by Duncan Thomas, 1990.

So Thomas is going to be an example of looking at mothers' and fathers' income on child health. OK? So you can interpret this in two ways.

Either moms care more about kids, right? That would be the ex ante way. Maybe moms care what more kids. We're going to say that child's health is mom's good.

Or maybe we could just say, we don't actually a priority of what's a mom's good and what's a dad's good. But we can just look at differences in the q 's as a test of income pooling. OK. And he's going to focus on nonlabor income for the reason that we talked about, which is that at least he doesn't have to deal with the fact that the labor income is endogenous, that the labor leisure decision is itself a choice.

We don't want that in there. We want the unearned income. So he's going to run a regression, which is going to look like some outcome variable as a function of the woman's unearned income and the man's unearned income.

And I think we talked about some of the challenges with this regression. So if I were to show you this, what do you think of this? So we talked with marriage market issues, matching issues. What else might you be worried about in this example? Yeah, Ahmed.

AUDIENCE: So is this a way of trying to get at the logging parameter in bold? I'm assuming the child is a common good for the marriage?

BEN OLKEN: This is a test of the pooled income hypothesis, OK? So this is a test which basically says-- so yes, in two respects. OK, sorry.

Yes, it is a test of that. The strict test of this is it's a test of the prediction that-- or I'll just show it up on the slides. It's a test of this prediction that fixing the bargaining weights, the identity of the income doesn't matter, OK?

So it's a test of that, right? It's testing, do we get different responses from Y_m or Y_f ? Y_f or Y_m ?

And yes, number 2, in terms of the outside option, what would be the alternative hypothesis? Yes, it's this idea here over here, exactly, that my outside option is going to depend on how much of the income is mine.

AUDIENCE: I'm not sure what you're trying to say. But the conceptual problem I'm having is that if it's true to model, still I'm thinking that bargaining-- given the differential bargaining power-- say it's fixed, it's exogenous-- I don't know-- 0.8, 0.2. It still feels like, for some reason, the income to the agent with the higher bargaining power could have a different effect on the common good. Are we assuming common good is equally valued by the agents?

BEN OLKEN: Yeah-- no.

AUDIENCE: They don't?

BEN OLKEN: But OK. But this is a great question. Actually, I was just thinking about this. This is a great question.

OK, so let's go back to this very big setup, OK? So in this setup, what happens if I change the fraction of Y -- if I change the fraction of unearned income that goes from the man to the woman?

AUDIENCE: For this model--

BEN OLKEN: Huh?

AUDIENCE: But it still doesn't change. Nothing--

BEN OLKEN: Yeah. So if the bargaining weights are fixed, so it's not in a world where they're determined endogenously like that-- if the bargaining weights are fixed, then nothing happens, right? That's the test they want to reject. OK?

And then we can reject that. One way we can reject it is the bargaining weights changed. I should have actually stopped and paused for other questions. Are there other ways in which we could reject this besides the bargaining weights?

I guess at some level, it's all the bargaining weight changing. But you seem to have something else in mind for some other way you could reject this besides the bargaining weights changing.

AUDIENCE: It seems like budgets can also be bargained over. You're bargaining over your [INAUDIBLE] if you don't have the budget in this model--

BEN OLKEN: What do you mean, bargaining over budgets?

AUDIENCE: The bargaining power can also determine-- say they have three different consumptions-- individual consumption, marriage consumption. My bargaining power in the marriage can determine, say, how much of the income of the other person I can get for the couple good and for myself as well. Say I don't care about my child. I have more bargaining power.

BEN OLKEN: But those are all-- it's all just a vector of consumption goods. It's x_1 to x -whatever.

AUDIENCE: I have to think more about it.

BEN OLKEN: Yeah, I don't think it matters. Another thing, by the way, which-- this is one microfoundation, by the way, for why it might fail. Another microfoundation for why it might fail might have to do with control rights. Maybe if we actually think about the mechanics of decision-making, it might be that actually, maybe we're not bargaining like this.

But maybe whoever has the cash actually makes decisions and presents the spouse with a fait accompli, right? We've all heard-- I feel like literature is filled with examples of this thing happening, as is real life. And that would be a slightly different microfoundation. I don't think it would-- that also you could think about ways in which this model could fail. Yeah?

AUDIENCE: Is the change in control rights model change in om or of?

BEN OLKEN: What, sorry?

AUDIENCE: Is the change in control rights different than our change in parameters in the bargaining model you covered [INAUDIBLE]?

BEN OLKEN: Yes. So in this example, so I haven't written down the change in the control rights version of this.

AUDIENCE: Wages-- the women's wages are deposited into a bank account that is managed by the man, and she knew that if her outside option-- that would similarly be the case in another marriage--

BEN OLKEN: Then it's going to be the same.

AUDIENCE: Right, yeah.

BEN OLKEN: Yeah. So I'm trying to think through exactly how you would write that down. It could end up looking similar, right? Because basically, in the outside option, I have control over the money. I just do whatever I'm going to want. So maybe it actually is going to look very similar, actually.

AUDIENCE: Yeah, it's the definition of control, right?

BEN OLKEN: Yeah, right? It's basically, if we don't agree, I just get to-- I think it's very similar. If we disagree, then I get to consume my stuff. And I get to decide how to spend my money.

You get to decide how to spend your money. And those are our outside options. So actually, sometimes, I think actually, it's a slightly different microfoundation in the sense of, in this one, in the example I gave before, I was thinking of the marriage splits up.

But the other one is, we don't agree on how to spend the money. And I just go spend my money. And you spend your money, right? That's another example. But it's going to look very similar in the formula in that sense. That's all I'm going to say.

OK. OK. Other challenges with this regression? Other things you want to think about?

OK, right. So we talked about-- there could be individual-level minimum variables. So for example, another thing about this regression is the utilities may be related. And this is a cross-sectional regression. So the utilities may be related to these income shocks, right?

So imagine that someone who grows up really rich and they're expecting to inherit a lot of money or whatever-- they may have different preferences than someone who doesn't grow up rich. And that might affect their consumption choices, right? So the ideal examples will have better variation than this, in just cross-sectional because you need to hold the utility-- the right test is we're holding the utility function constant and just shocking the Y's. Yeah, Patrick?

AUDIENCE: So in this case, do you have a shock in the fraction of shared income from each party? Or do you have a shock in total income, too?

BEN OLKEN: This one, they don't really have a shock. So the first version of this, they don't have a shock. It's just cross-sectional. I'll show you a shock in a second, OK?

So this one, they don't have a shock, right? So this is just going to be cross-sectional. And that's what I was saying.

The problem is of endogenous marriage formation. And the problem's that the utilities may be correlated with the income shocks. Right? So I was going to start by saying yes.

This was the first paper that did this in this form-- at least in development that I know of. He does find that basically, there are different income effects. So he looks at the ratio of the income effects for man's income versus woman's income and finds very, very substantial differences in those incomes.

But as you were saying, what you'd like to do is you'd like to have better identified shocks to the different income that are not going to be correlated with the utility shocks, with the preference parameters, OK? So let me not go through the results of this in detail. Let me go to the next paper, which was actually a paper by Esther, which she wrote in 2003, which looks at grandparents spending on their grandchildren.

And the reason for this, I think, not that grandparents and grandchildren are particularly of interest-- although, they of course, are interesting-- you could have equally written this paper with moms and fathers and their kids. But she had a policy change she could look at, which was a large unexpected change in people's unearned incomes. OK?

So in particular, at the end of apartheid, South Africa expands their pension programs to include Blacks. And there are a lot of kids-- 17% of kids who are living in multigenerational households that have a pension earner, OK? And she compares living with children who are living with an eligible grandparent-- so age 60 for females, or age 65 for males-- that's the eligibility rule-- to people who are living with a grandparent, but the grandparent is just younger, right? Yeah, Patrick?

AUDIENCE: The comparison takes place before they actually get the money, right?

BEN OLKEN: I thought it takes place after.

AUDIENCE: Wouldn't you be concerned about income effects?

BEN OLKEN: Right. Correct. So she's going to compare what happens if the grandmother gets the pension versus the grandfather gets the pension.

AUDIENCE: OK. OK.

BEN OLKEN: So it's going to be [INAUDIBLE], sorry.

AUDIENCE: If they find out they were going to be earning this income, could you look at before they got the actual--

BEN OLKEN: That would also be cool, as an anticipated thing. But I think the key thing is going to be comparing just-- so she wants to compare what is the treatment effect-- so there's two things going on in this regression, right? You have, say, 60-year-old female grandmothers and 59-year-old grandmothers.

And you have 65-year-old grandfathers and 64-year-old grandfathers. So we want to restrict ourselves to kids who are living with a grandparent. Because obviously, kids who are not living with a grandparent-- those are in different kinds of households.

The treatment effect of living with a grandmother is going to-- she's going to get that by comparing 65 or 64-year-old grandmothers, just eligible versus just not eligible-- kids living with. And the treatment effect of living with a grandfather is 65-year-old grandfathers versus 64-year-old grandfathers, right? So those are the two different treatment effects.

Both of those capture the income effect. Then the thing she wants to look at from the perspective of this topic is does it matter if the money goes to the grandmother or the grandfather? That's the kind of unitary-- that's the decision points here.

OK, question. First of all, does that seem sensible? And does that map into the framework I was just talking about? Yeah. Remind me your name, sorry.

AUDIENCE: Erin.

BEN OLKEN: Erin.

AUDIENCE: Yeah, I guess the one threat could be that sometimes, norms around whether or not an elderly grandparent can come back and live with a family might depend on what their gender is.

BEN OLKEN: The what, sorry?

AUDIENCE: Norms around whether a grandparent can come to return and live with the family, for their children, depends on the gender of the grandparent. I know at least in India, we were seeing that-- that's something we did with elderly females, who were more likely to live alone.

BEN OLKEN: Exactly. Right. So what we'd like to do-- so I think the cleanest test-- I was just thinking about this. I think the cleanest test of this would be kids who are living with both grandparents, right? So we're going to hold all of those household composition things the same.

And we'd exogenously vary which grandparent is just eligible. And she includes a lot of demographics. And I think this is in there, but we can double-check.

The cleanest test is households with grandparents of both genders comparing the case when the grandfather is just eligible to the case where the grandmother is just eligible. That would be the cleanest test. Yeah.

AUDIENCE: And the concern with that would be that there might be some specific dynamic in their relationship if the grandmother is older than the grandfather, potentially. Well

BEN OLKEN: They, don't have to be because of the differential age rule. But yes, it could be an issue. But I'm saying--

AUDIENCE: The fact that the grandmother became eligible before the grandfather implies that they're in the type of marriage where the woman would be older.

BEN OLKEN: No. Because the age difference-- the eligibility. The women are eligible five years younger.

AUDIENCE: I see.

BEN OLKEN: Yeah. And by the way, you are right to think that there could be discontinuities and that sort of thing. There's a paper in the US, for example, about earnings of spouses that seems to suggest that there is an important discontinuity at 50/50. So there could be a discontinuity exactly older or exactly younger. So you're right to think about that. But it's not a context-- not a concern here.

AUDIENCE: It would take five years--

BEN OLKEN: What?

AUDIENCE: So fewer than five years is enough of a buffer zone so that--

BEN OLKEN: Five years is probably not, I think, a particularly salient--

AUDIENCE: Discontinuity.

BEN OLKEN: Discontinuity. Right. It technically is, but I think in practice, it probably is not a salient one. Yeah.

OK. All right. So yeah. So this is what she's doing.

She's going to look at, in particular, she's going to control for these household composition things for the reason we talked about [INAUDIBLE]. But yes, that's the idea. So what she finds-- so by the way, so why is this better than the-- why is this an advance on the previous literature? It's because now we have a plausibly exogenous shock to income, right?

And so she's going to look at the weight for height for the kids as a measure of kids' nutrition. And she's going to look at it for girls and boys. And so what she's going to find is that basically-- so she's going to compare households where the woman is eligible to the one case where the man's eligible. And what she finds it she finds impacts of a woman being eligible on girls' nutrition.

She finds this impact of women being eligible on girls' nutrition and no impact of a man being eligible on girls' nutrition. She actually finds-- although this is not really the focus of the paper-- I don't know if she tests this-- in some specifications, it looks like the woman versus the man being eligible may also have near them as individually significant for the boys. Although I wonder if you test the difference of these coefficients, that the difference between them might be significant.

I'm not actually sure. Maybe not, actually, given the size those standard errors. But certainly, she finds this differential effect of men versus women on girls' nutrition.

OK, so that also suggests that-- so there's something going on where who you give the income to is affecting the patterns of consumption in the household. Yeah?

AUDIENCE: The difference of the coefficients being significant would mean what [INAUDIBLE]?

BEN OLKEN: That's a test. So the key test-- so going back to, I think, what somebody said that we don't want just income effects of someone being eligible. We want to test whether there's a difference of whether the man's eligible versus the woman's eligible.

AUDIENCE: Oh. So it's odd that that wasn't included in the table.

BEN OLKEN: Yes. Maybe it's in the paper and I missed it, but yes. If I was presenting this table, I would include P-value tests-- man coefficient equals woman coefficient.

AUDIENCE: So each coefficient-- it's a discontinuity in the [INAUDIBLE], correct?

BEN OLKEN: Huh?

AUDIENCE: So the coefficient that women eligible is like a discontinuity in age and something. Is that correct?

BEN OLKEN: OK. So--

AUDIENCE: Tied together, this is like a straight line, RD or something like that?

BEN OLKEN: OK. So I don't think this paper is written as an RD in the sense of-- we're going to talk about RDs a little bit later. Actually, I think the plan is for you to talk about RDs on Friday in recitation. As far as I remember-- and you can correct me if I'm wrong.

You can go check and correct me if I'm wrong. Or if Esther was here, she could correct me if I was wrong. You can ask her on Wednesday.

I don't think this paper is written as an RD with a continuous running variable in age. I think she doesn't have enough observations, for example, to do that, right? So there's only 1,500 observations here. So I think it's just dummies for being eligible.

The thought experiment is 100% an RD. But I don't think this paper is set up as an RD. But it could have been with more data. Other questions?

But yes. If you had infinite data and you were writing this paper today, I think you would set this up as a running variable. I think if I was writing the paper today, I would set this up-- and I had infinite data-- I would restrict myself to households with both a grandfather and a grandmother. And I would put in running variables for near the cutoff of age.

I would put in linear running variables for being on both sides of the discontinuity-- for birthday, for example. And I would have a discontinuous break in eligibility at the age cutoffs. And maybe that would be actually a good example to talk about on Friday, for example.

So is everyone following what I just said about regression discontinuity? Has anyone seen this before in at least one other context? Yes? Yes, OK.

Great. We'll go through it. We'll still go through it more detail in recitation on Friday.

But that's what I mean. OK. So given this finding, you might say, well, maybe policymakers may want to think about-- there might be policy implications with this in terms of giving resources differentially to men versus women. For example, there are lots of transfer programs in the world, which give resources to households to help them with a variety of things.

I'm going to talk about a bunch of them in the PF section. If you believe A, if you reject the unitary model, and B, if you believe some of these findings that women maybe, on average, channel more of their resources to their kids, and C, if you're a paternalistic decision maker or policymaker whose goal is actually to help the kids in the household, which you could be for reasons that we'll talk about in a few lectures, then all that suggests you might care about whether you give the transfer to the man or the woman. OK?

So here's an example of a recent randomized trial that looked at this. So this is in the context of the GiveDirectly trial, which was a trial that gave transfers to households. This is going to be a \$400 one-time transfer to households in Kenya.

And they did a bunch of things, some of which I'll talk about later. But one of the things they did is they randomized whether the transfer went to the men or the women in the household, in the case of a married household with a man and a woman, OK? So as I said, there's a bunch of stuff.

These are different outcome variables in the column. And this is the additional effect of randomizing who the recipient is to the woman. And in general, they get a joint P-value of 0.11 on a lot of things, like food security, health, education. They see nothing on a bunch of these.

Now they see nothing on a bunch of these. They do see some-- they have a psychological well-being index, which I think is a measure in a survey. And they have some questions on female empowerment.

Both of those seem to go up if we give the transfer to the women. But the other stuff, like health and education, don't seem to change. Now it is also worth noting that these don't seem to change from the treatment effect, either, of giving this transfer.

So maybe this was a case where the transfer didn't have differential effects on these incomes to begin with. And maybe it's therefore-- maybe it's not affecting these incomes to begin with. So maybe it's not so surprising that it had differential effects-- maybe.

But certainly, if you had just seen the previous papers-- the Thomas paper that says that the income going to the woman in the household yields substantially better nutritional outcomes for the kids, and likewise, the grandmothers and grandfathers paper-- a reasonable hypothesis might have been that even if you don't get much on average, or when you give it to the men, you get big effects if you give it to the woman-- they're not seeing that in this paper.

So I think it's also just worth pointing out-- it could be that this program didn't have effects on health and education for a variety of reasons. It could just be a different example. I think more work could be done on this question. Yeah, Patrick.

AUDIENCE: They talk about take-up rates between females?

BEN OLKEN: What's, sorry?

AUDIENCE: Take-up rates between males and females?

BEN OLKEN: Yeah. Oh, do they? I think the take-up rate is 100% in this example. I think this is all conditional on-- I don't remember this, but I don't think the take-up rates are an issue in this context. I could be wrong, but I don't remember that being an issue.

AUDIENCE: So everyone that got offered took it?

BEN OLKEN: OK, I don't remember for a fact. I don't remember low take-up as an issue in this paper. But you're right, sure.

If the take-up rate was different-- but you're not seeing a substantially lower-- on the did I spend the money variables, none of these look any different depending on whether they give it to the woman or the man. So if you thought the woman's take-it was a lot lower, you'd expect negative coefficients here. This is the additional effect of giving it to the women or the men. Yeah?

AUDIENCE: Could there be some sort of control thing at play? If you were to-- so in the previous example, I feel like part of the difference is that you're announcing the policy. And even if it's the case that women are just becoming eligible before the men would become eligible, the men can expect that they'll receive the 70% bonus at the time of becoming eligible. So--

BEN OLKEN: I'm sorry. Wait, I'm sorry. Say it again, sorry? Oh, you went back to the previous examples. Sorry.

AUDIENCE: In the previous example--

BEN OLKEN: Yes.

AUDIENCE: --you had an effect where you were announcing that you were going to give a transfer to both parties.

BEN OLKEN: Yes.

AUDIENCE: The transfer arrived to the woman faster than it did to the man in the case where she became eligible--

BEN OLKEN: Correct.

AUDIENCE: And so there's a notion that if the man waits sufficiently long, he'll get what's owed to him?

BEN OLKEN: Yes.

AUDIENCE: But in this particular example, you have this exogenous body that's just throwing money at a household. Could you make an argument that--

BEN OLKEN: Oh.

AUDIENCE: --there's a different control thing at play of this money, if it's being given totally randomly to one party or the other, is going to be pooled or is going to be put in the hands of one person, whereas that would less be the case?

BEN OLKEN: Yes. Yes. OK, sorry. Yes, I agree with you.

So totally plausible. So just to be clear, the argument was in the pension example, we're all going to get a pension someday. So maybe a reasonable bargaining solution is you do what you want with your pension.

I'll do what I want with my pension. It's all going to work out roughly because we're all going to get pensions eventually, whereas in this one, it's a one-time thing. And only one of us got it. And therefore, we really have to duke this one out. Is that basically your argument?

AUDIENCE: Yes.

BEN OLKEN: Yeah, it could be. That's interesting. Yeah. Yeah?

AUDIENCE: One [INAUDIBLE] this clearly not in this section that the transfer is for female or male rather than for the whole household when they do the experiment, when you give the transfer--

BEN OLKEN: Yeah?

AUDIENCE: Is it possible that the female or the male because they're already household, they think it is [INAUDIBLE]. So they are more likely to pool money together because it is not something really occurring, right?

BEN OLKEN: I think it's a-- yeah. I think that's a related context. What is the framing, in some sense? And how are they setting this up, right? Is that a reasonable--

AUDIENCE: Yeah, I'm just curious why there-- although it's given to female--

BEN OLKEN: So this was actually-- I think the way it was done is it was done through mobile money. And so the question is literally, whose phone does it go to? Your account or your husband's account?

AUDIENCE: I'm thinking like if I-- is it my earned income or is it my pension? It is clear it is for me. But for this case, if someone drops off money, and they met [INAUDIBLE] drop the money.

BEN OLKEN: Oh, I see. Yeah, yeah. No, exactly. It could be framed differently. That's what you're saying. It could be framed differently. I think what you're basically saying is-- so in the simplest bargaining model I set up, it's not different, right?

It's clear. Think about the version of the bargaining model I talked about where basically, it's like, I get to just make the consumption decision myself. If we don't agree, I guess we each get to spend our own money, right?

In that case, it's different because it went on my mobile money account versus your mobile money account, right? So in that version, I think it's still very different. I do think you are right.

And both of these comments I think are related to, what is the framing of this stuff and how this is presented to the households could matter in terms of what the implications are. That being said, I do think that we should update from this in the sense that you might have seen the previous results and drawn the totally reasonable policy inference that we should give the money to the moms if we care about the kids. And this says, actually, no, maybe it doesn't actually matter as much as you think in this example.

No, you disagree? Oh. It's very hard to read people's faces with the masks on. I'm doing my best.

AUDIENCE: Would an interesting comparison test be if you took the same income and then delivered it to both parents individually, and then slightly changed the proportion that you awarded to both parents, assuming you had many, many, many?

BEN OLKEN: Yeah, that would be a good test of your theory that basically-- that might be a good test of it. But again, it depends on the framing. And there are a variety of other-- I think Esther will talk about more of these next time. There are a variety of other more subtle tests.

We give you slightly more of something. We change the decision rights or we change the way we talk about it, or explicitly going into the details of this bargaining question. There is a bunch of additional work on that, which I'm not going to talk about today, but which I think Esther will talk about at least some of that next time.

OK. So now I want to get to the paper you guys read for today, which is about Pareto efficiency. So in the Burkina Faso example that you got-- this is another test of these models, right?

It's do we get to the efficient outcome? In this example over here, the efficient outcome is somehow we managed to earn S and then share it as opposed to each getting our own individual [INAUDIBLE]. In the Burkina Faso example that you guys read for today, individuals control different plots.

And so the Pareto optimal behavior is going to be to maximize total income and then share it according to some weights. OK? At least, that's going to be the benchmark. And we should talk about whether or not we think that's a reasonable benchmark.

Are there are other benchmarks you should think about as well? OK? And then the alternative is if I have control over the income from my plot and you have the control over income from your plot, we might not efficiently share resources across the two. And I have control over some resources.

And I might say, well, I'm going to overinvest in my own plot because I want the resources from my plot even if it was more efficient to put some of the resources in someone else's plot. So that's the question is are we jointly maximizing the production? Or somehow, is the identity of who controls the income from the plot affect the production choices on the plot? That's the basic question, right?

OK. So what's the setup? We have labor supply. We have consumption decisions. We have an area of each plot.

We have a production function, right? And we're going to assume this production function is going to be the same for all the individuals, although he's going to test that later in the paper, right? This is, I think, a key assumption, by the way. And we can talk about that.

I don't know what you guys thought of that. To me, when I was rereading the paper this morning, to me, a lot of my questions had to do with, is that assumption reasonable or not? We have total labor income, total labor supply.

And we have some joint decision making over consumption. So the Pareto efficient allocation is like we had before, right? We're going to maximize some weighted sum of the utilities subject to the budget constraint and the production functions.

And the key point-- and I'm going to come back to this idea-- the key point is that production decisions are independent of preferences. And so I'm going to actually talk about this in more detail again when I talk about the labor lectures, this idea that our production and our consumption decisions should be separable. It's a key test of a bunch of different models.

But basically, the idea is we're going to solve the-- we're going to maximize the total output by optimally allocating our resources across these two different plots to maximize output. And really, the key idea is it shouldn't matter for the household whose plot is the man's plot and whose plot is a woman's plot. We should just be maximizing output over the two plots. And then we should get the income and split it up.

OK. So that's basically the framework. He's going to define the yield on a particular plot as the amount of output per area planted. And he'll look at the deviation on some particular plot k -- or particular plot i , rather.

Crop k is going to be the-- they're going to look at the deviation from the output per acre on that particular plot compared to the average across all their plots. OK? And they're going to do this within crop, right? So they're going to say, we plant six different crops.

But actually, I plant peanuts. And you plant peanuts. And I plant rice. And you plant rice.

And we're going to look at my rice versus your rice is going to be the key we want to look at. OK? And so then they're basically going to run this regression, where they're going to run the output on whether it's the women's plot putting in household times time times crop fix effects.

So in this year, we produced this. We both produced rice. We both produced groundnuts, whatever. And we both-- it was a good year, or a bad year, or whatever. But we're going to look across our different plots.

OK? And the key test is going to be whether the output of the man's plot is different than the woman's plot. Yeah.

AUDIENCE: So on the previous slide, when you said, why not use between the [INAUDIBLE].

BEN OLKEN: Yeah.

AUDIENCE: [INAUDIBLE]

BEN OLKEN: Oh, why not?

AUDIENCE: Well if you measure g in terms of-- instead of yield and things like dollar rewards of yield--

BEN OLKEN: Oh, yeah. So yeah. So that's a good question. So why not do it all per dollar?

That's a good question, actually. Do you have an answer?

AUDIENCE: Well, I was wondering if there was a-- you don't want to overstrip the soil and plant too much corn. So you'd vary it by planting--

BEN OLKEN: In different places?

AUDIENCE: --sometimes, different-- yeah.

BEN OLKEN: Yeah, actually, that's a good question. I was trying to think that through. Why not do it? Yeah, go ahead.

AUDIENCE: I think it's between [INAUDIBLE].

BEN OLKEN: What, sorry?

AUDIENCE: Maybe they consume a lot of different crops to feed their home family rather than [INAUDIBLE] at market so that they really want to diversify the crop type. So they really put their decision-making if they want to get his share of his crop, his share of that crop so that they try to see, well, you switch fast and make a certain type of crop.

BEN OLKEN: Yeah. Yeah, no. The more I think about it, actually, I'm having a harder time answering it than I thought. So I was thinking, well, you want to be able to see about fixed effects, like for example, it's a bad year for this kind of crop, or a bad year for this kind of crop.

But why would one of them be systematically correlated with the woman's crop, for example? So why would we have the woman differentially plant? Well, one thing could be, for example-- no, I don't have a good example, actually.

Sorry, I don't have a good answer for you, actually. Someone else have a good answer? I think I see your point.

You should be able to dollarize it. It does seem like you're comparing more like to like in this example. Under the model, it's a good question.

Under the model, I can see what you mean. Under the model, we should have no differences in average-- everything should be equalized up to the margin. We should end-- well, if we had marginals, so we don't have marginals. We have averages.

So one reason-- maybe that could be related to it. But if we had marginals, you're right. I think we should be equalizing our output to the point where the marginal return in dollar terms is the same on every plot. So I think if we had marginals, I think that would be right.

It certainly can't hurt to do this. But I don't see a good reason not to do that. Yeah.

AUDIENCE: Why could we [INAUDIBLE]? I was going to ask a similar question. Wherever you do the misallocation, it seems like marginal or [INAUDIBLE] capital and the right people [INAUDIBLE].

BEN OLKEN: Well, sorry. That's the goal.

AUDIENCE: Yeah, yeah. So sometimes I think, what cases generally can we use averages to [INAUDIBLE]? Because it seems like that's what's happening here.

BEN OLKEN: Yeah, OK. So that's another good point. So OK, a couple of points.

Number 1-- so in the misallocation literature, they also end up looking at averages more often than you may think. They write down a theory, which is about marginals, right? But actually, they have to assume a functional form in order to-- what they observe is just output.

So you have to write down a functional form to get back the marginals from the averages. So I think a general point that I think all of you are making, which is totally right, which you're making, which is totally right is both in that literature and in this literature, the thing we really want is to estimate the marginals and randomly shock different things to estimate the marginal returns.

So the misallocation literature does that in functional form. This paper, I think, also essentially deals with it in functional form, right? So what are they going to do?

They're going to do two things, I think. What is he going to do? He's going to do two things to answer that.

The first is he's going to estimate these production functions, right? And so they look similar on the two different things. And that's, I think, going to-- and unless there's some big, unobserved term that would change the marginals in some way-- the fact that production functions look similar may help you.

The second thing, though, is he's going to directly look at inputs. And he's going to make this argument that the inputs look way way, way less. And the returns to the inputs, at least estimated cross-sectionally, look still high and concave. They look concave.

So he's going to basically say, if we think that the returns to these things are concave and we see much lower inputs on one versus the other, it is probably the case that the marginals are not being equated. That's basically the argument he's going to make. Is that a reasonable answer, or is that the answer to your question?

AUDIENCE: Yeah. When I asked the question, I was--

BEN OLKEN: What, sorry?

AUDIENCE: I just had a quick question on that.

BEN OLKEN: Yeah. So I think that's what they're going to do. But I think you're absolutely right to ask that question because I think lots of these papers make arguments about equating things which are true in marginals. But unless you have random shocks of things, you don't necessarily actually estimate the marginals.

Yeah. Other questions? OK. So the key feature in this context is that men and women control different plots, all right?

By the way, and he has data. As at least one of you pointed out in the comments, the data here is incredibly rich. To both have multiple plots and to have data on all the different plots is kind of really pretty remarkable.

OK? And he's going to focus on-- half the households, the man and the wives are harvesting the same crop on different plots of the same year. So that allows him to put in these fixed effects, although as per our previous point, if you could do it in-- oh, here's another reason, by the way, is that to do the thing in dollars, you need profits.

So the marginal-- the thing you want to think about-- first of all, you need profits. And you need marginal profits. So the marginals, we already talked about. The other thing which is tricky is profits.

You have to track off input cost. And that's actually tricky. So it's much easier to measure output than profits. So put that in mind, too.

Because you also need the inputs-- the inputs and the input costs. So that's the other thing that makes it easier to do it in output. And so in some sense, what is the key result, right?

The key result is that in general, that there's a negative coefficient on output for female plots controlling very flexibly for plot size and other plot characteristics, right? And I think one thing another-- something else you guys picked up on, which I think is nice about this paper, is they're controlling for these things very flexibly and not parametrically, right?

We don't really have a strong view of how soil type 12 versus 37 effects output or whatever. So he's going to control for all that pretty flexibly. OK. So let me pause here.

So what did you guys think of this? Did you find it convincing, not convincing? We talked about the marginals.

We talked about a couple of other issues. What are some other issues? What are some other thoughts you guys had about this paper? And generally, it seemed like you guys more or less liked it, but did other people have other concerns from reading it?

Come on. You all sent in comments. So does anyone have any other concerns about it? Yeah, Patrick.

AUDIENCE: [INAUDIBLE] the control steps were-- fully absorbed the quality of the plot.

BEN OLKEN: Yeah. No, right? So what you have to believe is that the controls are-- what he has is-- did I bring this? Tsh-tsh-tsh-tsh-tsh-tsh-tsh-tsh-tsh-tsh. Did I bring this?

No, I didn't bring this slide. But what he does is he shows you that basically, the results look pretty similar with and with-- actually, I think they look-- I don't [INAUDIBLE] look. But he shows you the results with him without the controls, right?

And I think you'd have to believe that the results still go through with the really rich controls. And you'd have to believe that there's a lot of other unobservable stuff that's not captured by these controls. But yes, I think that's something you could think about.

It would have to be that the-- but the other thing is he also shows you these direct results on input use. So you'd have to believe not just that the land was worse, but the marginal returns to the input use were a lot lower, which could be, right? But that would be the thing you'd have to believe.

Oh, right. Sorry. That's the other point. Actually, based on the observables, the women have better quality plots.

So you'd have to believe that the women's plots have some higher quality, but somehow the marginals are super duper low. And that's why they're not-- and the marginal returns and inputs are super duper low. Other thoughts? Other comments? Yeah?

AUDIENCE: [INAUDIBLE] think about the travel cost [INAUDIBLE]. Managing one plot seems--

BEN OLKEN: What, sorry?

AUDIENCE: --less costly. Managing one plot might be less costly than managing two plots, although the market [INAUDIBLE].

BEN OLKEN: Yeah, but the question is why does that explain this stuff, where you're basically using less of other inputs on the woman's plot?

AUDIENCE: Yeah, the location of the plots might be farther away from the husband, for example. So it's more costly to travel there, to travel to the best plot rather than just towards the plots around the house to reach my disproportionality [INAUDIBLE].

BEN OLKEN: Yeah. So that could be. Does anyone remember if he controls for distance to the plot?

AUDIENCE: There's one version where he just said one's very close.

BEN OLKEN: OK, right. [INAUDIBLE]. Yeah?

AUDIENCE: Wait. I forgot. Can plots be reallocate-- switched, reallocated? Or are they--

BEN OLKEN: So remember, this came up in some of the comments.

AUDIENCE: You mean moving the household? Yeah

BEN OLKEN: Yeah.

AUDIENCE: I think he was saying that it was allocated by the type of marriage.

BEN OLKEN: Exactly.

AUDIENCE: So I think it was [INAUDIBLE].

BEN OLKEN: Yeah.

AUDIENCE: And do we know if you use the plot more, the quality goes down?

BEN OLKEN: That is true. So well, I think you need to at least lay them fallow for some period of time. But there can be this notion of overuse. But would that be a problem for this?

AUDIENCE: I guess I'm wondering if the notion of-- if you're measuring that there's higher quality at a given time slice of the woman's plot that's just being used less-- that might--

BEN OLKEN: Yeah, right. So but right. But you'd have to tell some funky story, where basically the year they happened to do the survey was the year that all the women decided to let their plots be fallow or whatever, right?

There's no reason to think that would be particularly true in the survey year. And so you'd think it should average out. Other thoughts?

OK, I'm almost out of time. But I think you guys had a bunch of interesting comments on this. So I would encourage you to-- I think there's lots of interesting things going on here.

The last thing I wanted to mention, doo-doo-doo-doo-doo, was I just wanted to mention two things before I conclude because I think there's a lot of stuff here that I won't be able to talk to because Esther has more gender stuff she's going to talk about in the last lecture. I wanted to mention two things.

Number 1 is I wanted to mention-- some of you mentioned this idea of randomly dropping income on people. This paper by Robinson, I think, is a little bit closer to what some of you were talking about vis-a-vis the GiveDirectly example. So he's going to have random-- both spouses have random chance of receiving smaller income drops.

And look at the web of these things are different. That's just another randomized example I wanted to mention. And he does-- so I just wanted to mention that one as one example.

And the final thing I also wanted to mention that I'm not going to talk about is dynamics. So everything we've talked about has been static, right? And I just wanted to mention there's a whole additional literature on-- or there is a bit of an additional literature on dynamics.

And I wanted to make sure to have an opportunity to display my favorite cartoon on this subject, which is about dynamic allocation between parents and kids. But more seriously, there is this other idea about-- this Jensen and Miller paper, which is about whether or not I would-- if parents and kids can't contract, whether or not I want to do things to my kids to make sure that later, they'll behave nicely to me when they're grown up.

And you can think of that as about intergenerational friction, where I might-- in particular, the idea of this Jensen and Miller paper is that I may differentially underinvest in my kids because I want them to stay on the farm and stay close to me and provide old age support for me. I will not invest in them or give them the ability to go off and move off to the city, which might maximize our total income as a collective multigenerational household, but might mean that I might get less of it.

And so I just wanted to flag that, that these issues are not just static. We've talked about them statically, about man versus woman's income, and how does that affect choices. But it could also be parents and kids. And understanding that this heterogeneity, this dynamic optimization frictions is another interesting thing we can talk about.

OK. Sorry there was a lot there. I didn't get to finish. But Esther will take it over next time, then I'll come back after that. OK, thanks.